

The Role of Speaker and Hearer in the Character of Demonstratives

JEFF SPEAKS

University of Notre Dame
jspeaks@nd.edu

Demonstratives have different semantic values relative to different contexts of utterance. But it is surprisingly difficult to describe the function from contexts to contents (i.e. the character) which determines the semantic value of a given use of a demonstrative. It is very natural to think that the intentions of the speaker should play a significant role here. The aim of this paper is to discuss a pair of problems that arise for views which give intentions this central role in explaining the characters of demonstratives. As will emerge, these problems lead quickly to a foundational question about the semantics of demonstratives and many other context-sensitive expressions: the question of whether, in explaining their characters, we need to understand them as sensitive, not just to facts about the psychology of the speaker of the context, but also to facts about the audience of the context. I critically examine Jeffrey King's theory of demonstratives, which answers this question in the affirmative, and argue that it ultimately collapses into a purely speaker-based view of the character of demonstratives. I then show how to develop a much simpler speaker-based theory which both handles all of the cases handled by King's theory and avoids some of the more serious problems which King's theory faces. Towards the end I consider how we might solve the very difficult problems which result from cases in which speakers use demonstratives with conflicting referential intentions.

1. Indexicals and intentions

Some expressions have different semantic contents relative to different contexts of utterance. To try to state the features of context which determine the semantic content of a given expression type in a context is to describe, in Kaplan's sense, the character of that expression type.

In the case of at least some pure indexicals—like 'I' or 'now'—this looks reasonably straightforward. Although there are puzzles even here, to a first approximation it is plausible that the semantic content of 'I' is the speaker of the context, and the semantic content of 'now' is the time of the context.¹ But not all cases of context-sensitivity are so

¹ For discussion of some of the problems with pure indexicals, see [Mount 2008b](#) and [Cohen and Michaelson 2013](#).

easily handled; in many cases, it is far from obvious how to describe the character of the relevant context-sensitive expression. Examples include simple and complex demonstratives, comparative adjectives, counterfactual conditionals, and quantifiers.

Not all of these expression types are uncontroversially context-sensitive. But what should be uncontroversial is that, if they are context-sensitive, their characters are not as simple as the characters of ‘I’ and ‘now’. Many, for example, take the semantic contents of comparative adjectives like ‘rich’ relative to a given context of utterance to include a reference class — a set of things relative to which the thing in question is being said to be rich. But what about the context determines the relevant reference class? Others take the semantic contents of quantificational expressions to include some restricted domain of quantification — but what about the context determines the relevant domain restriction?

In answering these sorts of questions, it is very natural to think that the intentions of the speaker should play a significant role. The relevant reference class for ‘rich’ in a context seems to have something to do with the class of things the speaker intended to compare the relevant object to; and the relevant domain restriction for a quantificational expression in a context seems to have something to do with the class of things the speaker intended to be making a general claim about.

The aim of this paper is to discuss a pair of problems that arise for views which give intentions this central role in explaining the characters of indexical expressions like those just mentioned. My focus will be on what is arguably the simplest of the above cases — the character of demonstratives — but the problems (and, I hope, the solutions) are general. As will emerge, these problems lead quickly to a foundational question about the semantics of context-sensitive expressions like those on the above list: the question of whether, in explaining their characters, we need to understand them as sensitive, not just to facts about the psychology of the speaker of the context, but also to facts about the audience of the context.²

² This sort of foundational question can be asked about a variety of sorts of context-sensitivity which go beyond the case of explaining the characters of expressions like those just listed. One example is the attempt to explain various linguistic phenomena in part in terms of the common ground of the conversation. The common ground of the conversation, in Stalnaker’s sense, is a property of a conversation which is determined, not just by the properties of the speaker, but also in part by the properties of members of the audience. I think, though I do not argue the point here, that many of the problems raised below for

2. Two problems for the Simple Intention Theory

A natural first pass at explaining the character of demonstratives is the following theory:

Simple Intention Theory:

The value of a use of a demonstrative in a context is *o* iff the speaker intends *o* to be the value of the demonstrative in that context³

This seems to handle many cases quite well, as can be brought out by considering the problems which beset the competing view that the semantic value of a demonstrative in a context is determined by the accompanying demonstration.

On a somewhat flat-footed version of this view, demonstrations are pointing gestures, and such a gesture determines an object as its reference iff that object is the object pointed at — roughly, the object met by a straight line which extends the finger or arm of the pointing gesture.⁴ One immediate problem is that this theory is insufficiently general, because there appear to be uses of demonstratives accompanied by no demonstration (e.g. saying ‘that guy’ with no accompanying pointing gesture after someone has made himself especially salient). Other times, we manage to refer to someone by pointing at something else. (King (2014b) gives the example of saying ‘he’ while pointing at the chair of an absent student to refer to the student.) Indeed, as Reimer (1992) notes, given that when we point at something we always point at many things (the person, his shirt, the button ...) the case in which we refer to something while pointing at something else is the rule rather than the exception.

audience-based accounts of the character of demonstratives also arise for explanations of linguistic phenomena in terms of the common ground of the conversation (as opposed to, for example, speaker beliefs about the common ground).

³ Here I take the explanandum of a theory of character to be an account of semantic value relative to contexts, rather than an account of the semantic value of utterances of demonstratives. Since I take utterances to realize abstract Kaplanian contexts, the distinction will not matter too much for present purposes, and in what follows I will occasionally simplify by talking about the semantic values of utterances rather than of demonstratives relative to contexts. Thanks to an anonymous referee for pressing this point.

⁴ Something like this theory is defended by McGinn (1981), though he regards it as only ‘an idealisation of actual linguistic practice.’ He says: ‘the referent of a token of “that F” is to be the first F to intersect the line projected from the pointing finger, i.e. the F at the place indicated — one might almost say geometrically — by the accompanying gesture’ (p. 163).

Hence it seems that being singled out by a demonstration is neither necessary nor sufficient for being the semantic value of the corresponding demonstrative expression. The Simple Intention Theory, by contrast, seems to handle all of these cases rather neatly; in each case, the semantic value of the demonstrative is the object which the speaker intended to be the value of the demonstrative. Further, the Simple Intention Theory, unlike this sort of demonstration theory, seems to generalize to the other sorts of indexical expressions mentioned above; while we do seem, in the case of comparative adjectives, to have intentions to pick out a certain reference class, it is a bit difficult to point at a reference class.⁵

The Simple Intention Theory, though, faces problems of its own. Before moving on to those which will be my main focus, I would like to briefly mention two sorts of problems about how the theory is to be understood.

The first concern is about the content of the intentions which the Simple Intention Theory requires speakers to have. As stated, it requires speakers to have intentions about objects being the semantic values of expression tokens. And while I have intentions of this sort, it is at least not obvious that my two-year-old daughter, who is a competent user of demonstratives, does.⁶

It is not altogether obvious how this awkwardness can be avoided. We might change the theory to say that what matters are not speaker intentions about semantic values, but simply speaker's intentions to speaker-refer to objects. But it is not clear that this gets around the problem, since it is not obvious that my daughter has (or needs to have) intentions about speaker reference in order to use demonstratives. And even if it did, it is not clear that this would always lead to the right results, since, given that speaker reference and semantic reference can diverge, it is not at all clear that intentions to speaker-refer should always fix the semantic value of a demonstrative.

This is a serious problem, not just for the Simple Intention Theory of demonstratives, but for attempts to use speaker intentions in describing

⁵ One might reasonably object that there are theoretical options here other than the demonstration theory and some sort of intention-based theory. Here I am simply using the problems with the demonstration theory to motivate an intention-based theory; my main question will be how, if a theory of the latter sort is true, it is best elaborated. Hence I leave open the possibility that an alternative to both the demonstrative and intention-based theories is in the end preferable. For such alternatives, see Gauker's (2008) account of demonstrative reference in terms of a set of 'accessibility criteria' and the salience-based account given in Mount 2008a.

⁶ Thanks to Kent Bach and an anonymous referee for helpful discussion here.

the character of comparative adjectives, expressions of quantification, and the other context-sensitive expressions mentioned above. But here I will set it to the side. Even if it is a bit unclear whether all competent users of demonstratives must have intentions with respect to the semantic values of those demonstratives, moderately sophisticated speakers sometimes do have intentions of this sort; the cases on which I focus will all be cases of this sort. The problems on which I will focus all remain even if we (admittedly, artificially) restrict our attention to cases in which speakers do have the right sorts of intentions.⁷

The second, less serious problem about how the theory is to be understood concerns the use of the phrase ‘semantic value’, not just in specifying the content of the relevant speaker intention, but in stating the target of the theory. The reason why I (following King 2014b, n. 3) say that we want an account of what determines the semantic value of a demonstrative in a context—rather than, for example, what determines its semantic content—is just to stay neutral between various views about the semantic content of demonstratives. The problems about how intentions determine the semantic content of a demonstrative in a context are independent of questions about whether we take demonstratives to be devices of direct reference or, for example, quantificational expressions.⁸ On the latter view, demonstratives will never have objects as their semantic contents; but we can still speak of objects as their semantic values, where the semantic value of a definite quantifier is the object (if any) which uniquely satisfies the relevant restrictor. In what follows I will occasionally talk about ‘reference’ and ‘reference failure’; these should be understood as shorthands for ‘semantic value’ and ‘failure to have a semantic value’.

With these clarifications out of the way, we can turn to the two problems for the Simple Intention Theory which will be our focus. The first is what I call the *problem of insufficient intentions*, and can be brought out by an example due to Jeffrey King (2014b, p. 224):⁹

The beach:

I am sitting on Venice beach on a crowded holiday looking south, with swarms of people in sight. I fix my attention on a woman in

⁷ I discuss this problem about the nature of the intentions posited by intention-based theories of character in more depth in *Speaks MS*. See also Gauker 2008, p. 362.

⁸ For views of the second sort, see King 2001 and Hawthorne and Manley 2012.

⁹ This scenario and the others which follow are presented in the words used by the cited author.

the distance and, intending to talk about her and gesturing vaguely to the south, say ‘She is athletic’. It seems quite implausible in such a case to say that I succeeded in securing the woman in question as the value of my demonstrative simply because I was perceiving her, intending to talk about her and gesturing vaguely in her direction.

Here we have a speaker intention of the sort required by the Simple Intention Theory, but, King thinks, no semantic value. Similar cases have been emphasized, and similar morals drawn, by Howard Wettstein, Marga Reimer, and Christopher Gauker.¹⁰

The problem of insufficient intentions is the problem posed by cases in which a speaker intends that some object be the value of her demonstrative, but, for some reason, that intention is just not enough to give that demonstrative a semantic value. In the cases we have considered, this seems to be due to the speaker’s failing to do enough to put her audience in a position to determine which object she intended to be the value of the demonstrative. If the judgements of King and others about these cases are correct, then the Simple Intention Theory is false.

In the second sort of problem case, we get a speaker intention of the relevant sort, but some semantic value other than the object singled out by the intention. Kaplan provides a well-known example of this sort (Kaplan 1970, p. 39):

Carnap and Agnew:

Suppose that without turning and looking I point to the place on my wall which has long been occupied by a picture of Rudolf

¹⁰ Gauker (2008, p. 363), for example, gives us the following case:

Sally is standing next to the mirror gazing toward the tie around Harry’s neck and says, ‘That matches your new jacket.’ As a matter of fact, Sally has been contemplating in thought the tie that Harry tried on two ties back. At first she thought she did not like it, but then it occurred to her that it would look good with Harry’s new jacket. We can even suppose that in saying ‘that’ what she intended to refer to was the tie two ties back. But under the circumstances, Harry is in no position to realize that the tie she intended to refer to was the tie two ties back and therefore is in no position to take Sally’s intention into account in identifying the reference of her demonstrative ‘that.’

In this case, Gauker thinks, for reasons parallel to the reasons King cites in the case of the beach, that the prediction of the Simple Intention Theory—that the demonstrative refers to the tie Harry was wearing ‘two ties back’—is simply mistaken. (Gauker adds to this the positive claim that the demonstrative does refer, contra the speaker’s intention, to the tie worn at the time of the utterance; this seems to me to be a mistake.) See also the version of the ‘Fido’ example in §IV of Reimer 1991. Wettstein (1984) suggests the general rule that ‘If the speaker fails to make his reference available, his speech act is defective, and not even the best intentions can repair the defect. The speaker, strictly speaking, has not asserted anything determinate, i.e., anything at all’ (p. 75). This, as we will see, is very close to King’s view.

Carnap and I say: [That] is a picture of one of the greatest philosophers of the twentieth century. But unbeknownst to me, someone has replaced my picture of Carnap with one of Spiro Agnew. ... I have said of a picture of Spiro Agnew that it pictures one of the greatest philosophers of the twentieth century. And my speech and demonstration suggest no other natural interpretation ... No matter how hard I intend Carnap's picture, I do not think it reasonable to call the content of my utterance true.

Another well-known case of the same general sort is due to Marga Reimer (1991, p. 90):

Suppose, for instance, that I suddenly realize that I have left my keys on the desk my (shared) office. I return to my office, where I find the desk occupied by my officemate. I then spot my keys, sitting there on the desk, alongside my officemate's keys. I then make a grab for my keys, saying just as I mistakenly grab my officemate's keys, 'These are mine'.¹¹

In the first case, the speaker intends to refer to the picture of Carnap, but does not succeed in making this picture the value of his demonstrative; in the second, the speaker intends to refer to her keys, but these are not the value of her use of the demonstrative 'these'. So these cases are *prima facie* counterexamples to the Simple Intention Theory.

It plausible that these cases pose a very different challenge than the cases used to exemplify the problem of insufficient intentions. In the case of the beach, we seem to have — though we will see reason for questioning this later — a case where we have an intention of the right sort, but for some reason it is just not enough to give the demonstrative any semantic value at all. The Carnap/Agnew case is different, because it is plausible — even if not obvious — that the demonstrative there *does* have a semantic value — namely, the picture of Agnew. One might think that this shows that the conditions given by the Simple Intention Theory are not necessary for reference — but (as others have noted) this neglects the fact that the speaker in that case, as well as intending to refer to a picture of Carnap, also intends to refer to the picture at which he is pointing. Analogous remarks apply to the example of the keys, and the intention to refer to the object toward which one's hand is moving. So here we have cases in which, intuitively, an intention fails to determine the semantic value of a

¹¹ See for further discussion of this case Bach 1992.

demonstrative not because it is just, in some way or other, ‘not enough’, but because it is overridden, or trumped, by another intention. Hence I will call this the *problem of conflicting intentions*.

Cases of insufficient intentions and conflicting intentions both challenge the sufficiency, but not the necessity, of the conditions given by the Simple Intention Theory. Hence it is natural to think that the problems with the Simple Intention Theory can be solved by strengthening the theory by adding one or more extra conditions. If the judgements about the semantic values of demonstratives in the above cases are correct, we would want these extra conditions to provide answers to the following two questions: (i) Why do some intentions of the relevant sort fail to give the demonstrative any semantic value? (ii) When we have two or more intentions of the relevant sort which single out different objects as the semantic value, why does one sometimes trump the other?

3. The Coordination Account

There are two broadly different ways of thinking about where we should look for these extra conditions. On the one hand, we might look for some extra mental state of the speaker—something other than the bare intention that *o* be the semantic value of the demonstrative—which we might require for reference. On the other hand, we might look, not to the speaker, but to her audience, and look for some facts about what the audience does or would do in response to the utterance which we might require for reference. As mentioned above, this is a very general choice point for theories of the semantics of indexical expressions; parallels of the problems of insufficient intentions and conflicting intentions can be raised for intention-based theories of the characters of comparative adjectives, quantificational expressions, and other plausible examples of context-sensitive expressions.¹²

Jeffrey King is an example of someone who thinks that we should strengthen the Simple Intention Theory by adding an extra condition which focuses, not on the psychology of the speaker, but on the audience. On this view, which he calls ‘the Coordination Account’,

the value of a given use of a demonstrative [in a context] is that object [*o*] that meets the following two conditions: 1) the speaker intends [*o*] to be the value; and 2) a competent, attentive, reasonable hearer would take [*o*]

¹² As others have noted. See especially Glanzberg MS and King 2014a.

to be the object that the speaker intends to be the value. (King 2014b, p. 226)¹³

This account has some notable virtues. One is that it seems to put demonstrations in exactly their proper place. Sometimes, demonstrations seem essential to successful uses of demonstratives; other times they seem entirely superfluous. The Coordination Account says why: sometimes demonstrations are essential in making clear to an audience a speaker's referential intentions; other times they are not. Crucially for our purposes, the Coordination Account also seems, unlike the Simple Intention Theory, to solve the problem of insufficient intentions and the problem of conflicting intentions.

The theory is more or less tailor-made to solve the problem of insufficient intentions. The intuition of reference failure in the cases described by King and others stems from the fact that, in these cases, the speaker has not done enough to put her audience in a position to discern her referential intention. The Coordination Account delivers the verdict of reference failure in these cases precisely by making it a condition on reference that speakers not fall short in this way.

Although King does not (for reasons which will be discussed below) apply the account in this way, one might also think that the Coordination Account can explain what is going on the case of Carnap and Agnew. Why does the intention to refer to the object at which the speaker is pointing determine the reference of Kaplan's demonstrative, while the intention to refer to a picture of Carnap does not? Simply because the speaker's audience would take the speaker to be intending to refer to the picture at which he is pointing, but would be in no position to figure out that the speaker intended a picture of Carnap to be the value of his demonstrative. Hence the picture at which Kaplan is pointing seems to satisfy both conditions of the Coordination Account, whereas the picture of Carnap satisfies only the first condition. If we agree with Kaplan that the picture of

¹³ An important predecessor of King's view is the view defended by Wettstein (1984): 'One who utters a demonstrative is responsible, from the point of view of the natural language institution, for making his intended reference available to his addressee, and so he is responsible for the cues that a competent and attentive addressee would take him to be exploiting. The cues for which he is responsible, those that he, to all appearances, exploits, are the cues that determine the reference. ... The gap between meaning and reference is to be bridged by the cues that the competent and attentive addressee will reasonably take the speaker to be exploiting' (pp. 72–3). Here I focus on King's version, since I think that it is the most well-developed version of this sort of audience-based view of character in the literature.

Agnew is the semantic value of the demonstrative, this seems like exactly what we should want.

Neat as these solutions seem to be, in what follows I will argue that they dissolve upon closer inspection. And the reasons why these explanations fall short, I think, will tell us something about the correct answer to our foundational question about whether the characters of demonstratives (and other context-sensitive expressions) should be sensitive to the properties of audiences as well as of speakers. In particular, I will argue, we will find strong reason to prefer a purely speaker-based theory of the character of demonstratives over a theory like King's.

4. Coordination and the problem of insufficient intentions

Let us call the second condition on reference that King gives the *coordination condition*.¹⁴ In this section, I want to focus on how, exactly, this condition is to be understood, if we want the Coordination Account to solve the problem of insufficient intentions. My aim will be to show that consideration of cases forces us to an understanding of the coordination condition on which it is not really a condition on audiences at all, but a further condition on the psychology of the speaker—and that, once we reach this point, we will be able to see that much simpler, and more plausible, conditions on the psychology of the speaker are available.

Let us stipulate, with King, that a competent, attentive, reasonable hearer would know, in Stalnaker's sense, the common ground of the conversation. That is, such a hearer would know which propositions are mutually believed to be accepted for purposes of the conversation (Stalnaker 2002). Then a first pass at understanding the coordination condition might be

(Coordination-A) [$\forall x$: x is a member of the audience of S 's utterance] (x is competent, attentive, etc., and knows the common ground of the conversation $\Box \rightarrow x$ would take o to be the object that S intends to be the semantic value of d in c)

So, on this interpretation of King's theory, an object will be the semantic value of a use of a demonstrative just in case that object both

¹⁴ King (2013, 2014b) considers a few different ways in which this condition might be understood. My way of laying out the options here differs a bit from his, but I think that every option he considers is equivalent to one of the interpretations discussed below.

satisfies the condition laid down by the Simple Intention Theory, and satisfies (Coordination-A).

Given the focus of the Coordination Account on the audience, it is worth remembering that demonstratives can be used in thought as well as in communication. This might lead us to wonder: can the Coordination Account handle examples of talking to oneself?

The answer is: yes, and very neatly. We can bring this out by considering an example King discusses:¹⁵

The idiot:

Lying in bed one night I say to myself, ‘That guy was an idiot’, remembering someone that I met at the bar. Given that I am only talking to myself, it seems as though I can succeed in giving ‘that guy’ a semantic value, despite (as in the example of the beach) not having done anything which would have made clear to a third party which idiot was the object of my referential intention.

It is natural in cases like this to identify the audience with the speaker. Plugging this into the Coordination Account, we get the result that ‘that guy’ will have a semantic value iff the speaker has put himself in a position to know his own referential intentions. But, except in odd cases in which we have a failure of self-knowledge, the speaker always will know his own referential intentions, and clause (2) of the Coordination Account will be satisfied. That means that, for these cases, the Coordination Account will be equivalent to the Simple Intention Theory. And this seems like intuitively exactly the right result. The sorts of cases which seem to be counterexamples to the Simple Intention Theory—like the utterance of ‘She is athletic’ in the case of the beach—do look like cases of reference failure in cases in which you are talking to someone else, but seem perfectly fine in the case where you are talking to yourself.

This is an especially nice result, because other sorts of theories that make use of audiences, or intentions directed at audiences—like Gricean analyses of speaker meaning, or certain views of conversational implicature—end up foundering on cases in which, intuitively, the speaker and the audience are identical.¹⁶

¹⁵ See King 2014b, n. 26. King attributes the example to Andrew Moon.

¹⁶ The problem for Gricean analyses of speaker-meaning, in a nutshell, is that one can mean something when talking to oneself without intending that one recognize one’s own mental states on the basis of the recognition of that very intention. But, if we simply identify speaker and audience in the case of talking to oneself, that is exactly what the Gricean analysis

But these are not the only ‘audienceless’ cases, and some of the others are not as easily handled by the Coordination Account. These are cases which involve mistakes on the part of the speaker about his audience. Here is one:

The sneaky students:

I teach Philosophy 101 in a large auditorium which darkens during the lecture so that the students can better see the slides; in fact, though, it becomes a bit darker than it needs to, to the point where I cannot see the students during the lecture. The students have figured this out, and now, very quietly, exit the room minutes after the lights go down, and return minutes before the lights go back up. In the interim, I’m speaking to an empty room.

During the lecture, I might use plenty of demonstratives; and it seems to me very clear that many of the sentences in which those demonstratives figure, including many simple predications, might be true. Hence it seems clear that many of the demonstratives I utter will have semantic values. If, pointing clearly and carefully at the lectern, I say ‘That lectern ...’ it seems clear that the semantic value of ‘that lectern’ is, just as it would be had the students not left, the lectern. But I have not done anything to make my referential intentions clear to my audience, there being no audience to make them clear to.

What is special about this case is not really that it is an audienceless case; it is rather that it is a case in which a speaker is mistaken about whether he has an audience. This blocks the move suggested above, on which we simply identify the audience with the speaker. It is not just that this seems intuitively in-ap in a case where the speaker takes himself to be addressing a large audience; it is also that it gives the wrong results in these cases. This can be brought out by considering a possible continuation of the story of the sneaky students:

The speckled hen:

During the lecture, I discuss the example of the speckled hen, and show a picture of a many-speckled hen on the screen. Carelessly gesturing toward the hen, intending to single out one of the speckles, I say ‘That speckle ...’

This seems to me to be relevantly like the example of the beach, which was used to argue against the Simple Intention Theory. I did not do

of speaker-meaning would require. See, among other places, Harman 1974, Chomsky 1975, and Speaks 2008.

anything like enough to make clear to my intended audience what the intended referent of the demonstrative was; so, given the intuitions that drive the Coordination Account, it seems clear that this should be a case of reference failure. On the version of the view which employs (Coordination-A), the Coordination Account thus has an insufficiently general solution to the problem of insufficient intentions.

One might dig in one's heels here, and say that these cases really are not the same—one might insist that 'that speckle', just like 'that idiot', has a semantic value simply in virtue of the speaker's knowledge of her own intentions. But this is very hard to believe. Suppose that we modify the example of the beach so that, just before your utterance of 'She is athletic', your interlocutors are quietly assassinated. Could this really be the thing that gives a semantic value to your utterance of 'she'? There seems to be no principled reason why reference failure of the sort which King takes to be exemplified by the example of the beach could not also occur in cases in which one mistakenly takes oneself to have an audience other than oneself.

Fortunately, it looks like there is an easy modification of (Coordination-A) available. We might plausibly go for something like

(Coordination-B) S's utterance has an audience $\Box \rightarrow [\forall x: x \text{ is a member of the audience of S's utterance}]$ (x is competent, attentive, etc., and knows the common ground of the conversation $\Box \rightarrow x$ would take o to be the object that S intends to be the semantic value of d in c)

This is equivalent to (Coordination-A) for cases in which S does have an audience, but appropriately distinguishes between the cases of the sneaky students and the speckled hen, since it seems true that were there an audience for my lecture, they would have been in a position to discern my referential intention in the case of 'that lectern', but not in the case of 'that speckle'.¹⁷

¹⁷ One might think that this leads to some rather odd results. (Coordination-B) improves on (Coordination-A) by, in effect, stipulating that we always evaluate reference at a world where we have an audience whose properties we can use to evaluate referential success or failure. But this leads to immediate trouble when we consider cases in which a speaker is demonstrating her audience, like the following:

The hallucination:

I am having a convincing hallucination that I am having a conversation with Alice, Bob, and Carol, and pointing at (what I take to be Carol), say 'She ...'

(Coordination-B) is not, however, a comfortable stopping point. As it stands (given the usual Stalnaker-Lewis semantics for counterfactuals or something similar), (Coordination-B) asks us to look at the world *w* most similar to the world of the context in which *S* has an attentive audience who knows the common ground of the conversation, and ask whether *S*'s audience in *w* would take *o* to be the object to which *S* intends to refer.¹⁸ The problem is that this will often give us the wrong result in cases in which the speaker has false beliefs, not about whether she has an audience, but rather about the properties of her audience. One type of such mistake comes when a speaker is confused about who the members of her audience are. We can imagine a variant of the case of the sneaky students which fits this description:

The replacement audience:

The students have grown worried that, despite the darkness of the classroom, I will notice that there are no bodies in the seats. So they hire a surprisingly well-behaved troupe of three-year-olds to sit quietly in their seats during class.

These three-year-olds, we can suppose, are perfectly attentive, and yet are in no position to understand what I intend to refer to when I use the demonstrative 'that lectern'. But, as in the original case of the sneaky students, this seems to me plainly not sufficient to stop that demonstrative from having a semantic value. As in that case, it suffices to note that it is extremely plausible that a simple predication involving the demonstrative might well be true.

Similar problems are posed by a case in which one takes oneself to be one's only audience, but someone else is listening in, as in the following continuation of the case of the idiot:

The timid burglar:

Unbeknownst to me, as I say 'That guy was an idiot', there is a burglar hiding beneath my bed. (I got home mid-burglary, and

I intend to refer to Carol, so clause (1) of the Coordination Account is satisfied. But (Coordination-B) is also satisfied, since, we can suppose, in the nearest world in which I have an audience, that audience would include Carol, and would hence would be in a position to figure out that Carol is the object of my referential intention. But that just seems wrong; in this context, my use of 'She' seems like a clear example of reference failure. Ultimately, this is probably best thought of as an instance of the problem of conflicting intentions, since it looks like a case in which the speaker's intention to refer to the person at whom he is pointing trumps his intention to refer to Carol. More on this sort of case below.

¹⁸ Here I am ignoring the possibility of multiple worlds equally similar to the world of the context; I do not think that anything hangs on this, for present purposes.

the burglar thought it best to escape detection by hiding under my bed until I went to sleep.) The burglar, having nothing else to do, is listening to my monologue, and wonders who the idiot in question is; he wishes that I had done enough to make my referential intentions clear to him, but I haven't.

Surely the presence of the burglar should not change our view that 'that idiot' had a semantic value, and that the sentence I utter is true iff the guy in question was, in fact, an idiot.

It is natural to respond to this case along the following lines:

The burglar did overhear your utterances, but that is not enough to make him, in the relevant sense, your audience; for that we need in addition that you believe that you are addressing him, a condition which is plainly not satisfied in the present case.

And, one might think, if we say this about the burglar, perhaps we could say something similar about the case of the replacement audience; could the absent students rather than the three-year-olds count as your audience despite their not being present for the relevant utterances? After all, you do think that you are speaking to them.

All of this is very plausible. But it is worth noting that this requires some account of what it is for someone to be, or not be, a member of a speaker's audience, which is presumably to be given in terms of the speaker's beliefs or intentions—and to make this move would be to move away from a theory given partly in terms of the properties of an audience toward one given wholly in terms of the psychological states of the speaker.

But I will not dwell on this sort of concern here, since problems parallel to those raised by the replacement audience and the timid burglar arise in cases in which the speaker knows who his audience is, but is mistaken about what one or more members of the audience are like, as in the following case:

Sudden blindness:

I'm having a beer with a friend at a bar, and, pointing to her glass, say 'That beer looks flat'. Unfortunately, she was struck blind moments before my utterance, and hence was unable to discern the object of my referential intention.

My friend's unfortunate condition plainly does not stop 'that beer' from having a semantic value—after all, it does not stop the sentence I utter from being true iff that beer really does look flat. But

(Coordination-B) implies that it does; after all, I do have an audience for my utterance, and she was (let us assume) perfectly attentive and competent. Hence the antecedent of the counterfactual is true in the world of the context, and the condition is satisfied iff the embedded counterfactual is true. But the antecedent of that counterfactual is also true in the world of the context. And if this is the case, condition (Coordination-B) is satisfied iff my friend actually discerns my referential intention. Since she does not, it is not.

It is worth pausing to say why the antecedent of the embedded counterfactual is true in the world of the context. To show this we need to show, first, that my friend might well be competent and attentive; and, second, that she might well know the common ground of the conversation. One might suggest, against the first point, that we build into the conditions abbreviated as ‘competent, attentive, etc.’ that the audience have, among other things, normal eyesight. This would solve the problem; but it would build too much into the idealization of the audience, since it would entail that I could successfully use demonstratives in the ordinary way to an audience which I know to be visually impaired — and this conclusion would be very much at odds with the intuitions which are supposed to push us away from the pure intention theory to the Coordination Account.

On to the second point: does my friend know the common ground of the conversation? One might think not, on the grounds that this conversational context is a defective one in which the speaker is radically mistaken about his audience; perhaps this is enough to deprive the context of any common ground at all. But this line of reasoning is a mistake so long as we are understanding ‘common ground’, as is standard, in Stalnaker’s sense:

It is common ground that ϕ in a group if all members accept (for the purpose of the conversation) that ϕ , and all believe that all accept that ϕ , and all believe that all believe that all accept that ϕ , etc. (Stalnaker 2002, p. 716)¹⁹

In the case of sudden blindness, plenty of propositions satisfy this condition — for example, the proposition that each participant is drinking a beer. And there is nothing to stop my friend from knowing that all of these propositions are commonly accepted, and hence knowing the common ground of the conversation.

¹⁹ King (2014b, p. 226) indicates that he has Stalnaker’s conception of common ground in mind.

Once we see the recipe, it is clear that less dramatic examples can also be used to provide counterexamples to (Coordination-B). A plausible thought is that what has gone wrong here is that (Coordination-B) leaves open the possibility that referential failure or success can turn on properties of my audience which are hidden from me and other conversational participants. What we want is a way around this, which must be a way of modifying (Coordination-B) in such a way that the antecedent of the embedded counterfactual takes us to a world in which my friend has normal vision.

Schematically, the natural way to do this is to define some properties of the audience—call them the X-properties—which are such that what matters for referential success is the question whether the audience discerns the speaker's referential intention in a world in which the audience is competent, attentive, knows the common ground, *and has the X-properties*. That is, we ought to look for some appropriate instance of

(Coordination-C) S's utterance has an audience $\square \rightarrow [\forall x: x \text{ is a member of the audience of S's utterance}] (x \text{ is competent, attentive, etc., and knows the common ground of the conversation and } \underline{\text{has the X-properties}} \square \rightarrow x \text{ would take } o \text{ to be the object that S intends to be the semantic value of } d \text{ in } c)$

The obvious next question is: what are the X-properties? Given that we are trying to handle cases like sudden blindness, it looks like these properties should have something to do with the speaker's—or the conversational participants'—conception of the audience.

A natural candidate—which King considers (though not using the 'X-properties' terminology) is that the X-properties should be the properties attributed to the audience by the common ground of the conversation. That is, roughly: every property F which is such that all conversational participants accept for purposes of conversation that the audience has F, and all believe that all accept this, and all believe that all believe that all accept this, etc.²⁰

Does this help with Sudden Blindness? Only if the common ground of that conversation attributes being normally sighted to my friend; and that will be true only if my friend accepts for purposes of conversation that she is normally sighted. And she may—even though

²⁰ King (2014b) considers this as an option; he endorses it in King 2013.

she of course does not believe that she is normally sighted.²¹ But she also may not—she may, for example, think that I have noticed the sudden change in her condition. The case might be elaborated either way—and, crucially, which way it is elaborated seems completely irrelevant to the question of whether ‘that beer’ has a reference. Either way we describe the case, it seems clear that the sentence uttered is true iff the beer in question really does look flat—and hence also clear that ‘that beer’ has a semantic value (namely, the allegedly flat-looking beer).

One might worry that this case is too far-out to be taken seriously as a test of a theory of demonstratives.²² But, again, once we see the recipe, it is not hard to come up with other less far-out examples. All we need is a case in which I am slightly mistaken about the common ground—a case in which I have false (but perhaps quite well justified) beliefs about what information is mutually believed to be accepted for purposes of the conversation. Imagine, for example, that I am sitting at a basketball game next to someone I do not know, and suppose that one of the teams has a 6’10” (2.1 metre) point guard. Without pointing, I might say ‘That guy’s really tall for a point guard’. Here I am assuming that the person sitting next to me knows enough about basketball to know that the point guard is the guy who dribbles the ball up the court, and that they usually are not very tall. These are quite reasonable assumptions to make about someone in the crowd at a basketball game. But he might well not know that—he might be at his first game. And that should not stop my use of ‘that guy’ from having a semantic value, and should not stop the sentence I utter from being true.

And there is a more general point here. We have already accepted that speakers’ justified but false beliefs about whether they have an audience (the sneaky students), who is in the audience (the replacement audience), and the properties of their audience (sudden blindness) should not deprive their words of a semantic value. It is hard to see, in principle, why justified but false speaker beliefs about the common ground should be any different.

²¹ Recall that, on Stalnaker’s account of common ground, a proposition can be a part of the common ground even if it is not mutually believed, so long as it is mutually believed that the proposition in question is mutually accepted for purposes of conversation.

²² King expresses this worry in King 2013.

This case—like sudden blindness—makes it seem like what matters is not what the common ground in fact ascribes to the audience, but rather what the speaker believes about the audience. And this might suggest a natural candidate for the X-properties: perhaps they are just the properties that the speaker believes his audience to have.²³

This looks like progress: it seems, for all we have said, to handle sudden blindness, and indeed all the cases discussed in this section. But it leads to a question. Our present understanding of the coordination condition defines referential success partly in terms of the properties the speaker takes every member of her audience to have. But which of these beliefs matter?

As stated, the answer is: all of them. We look at the nearest world in which *every* property which the speaker attributes to every member of his audience is possessed by every member of the audience. But it is immediately clear that we need some restriction on the class of relevant beliefs.

One way to bring this out is by noting that I can have beliefs about whether my audience will discern my referential intention, and that these beliefs can be mistaken. Suppose that we add to the example of the beach that I believed that my audience would discern my intention to refer to the particular woman to whom I in fact intended to refer. Then our present interpretation of (Coordination-C) will have us look at a world where this belief about my audience is true, which of course ensures that in that world my audience will discern the object of my referential intention, which in turn will entail, given the relevant version of the Coordination Account, that my use of the demonstrative is

²³ King (2014b, p. 228) also considers the view that we should consider only worlds in which the audience has the properties the speaker knows him or her to have. But that will not help with cases like Sudden Blindness, since I do not know that my audience is normally sighted—after all, she is not.

It is worth pausing for a moment to be clear about the way in which I am thinking of ‘the properties a speaker actually believes his audience to have.’ Let us call these properties ‘@B.’ A first try at defining the set of properties which are conjoined to form @B might be

F is a member of @B iff $\forall x$ (x is a member of the audience of S’s utterance $\rightarrow$ S believes that x is F)

But this will not work, for two reasons. (1) In cases in which the speaker has no audience, it will trivially entail that every property is part of @B. (2) In some cases, such as speeches to large crowds, there may be no particular individual such that S believes of that individual that she is speaking to him. This suggests that we should go for something more like:

F is a member of @B iff S believes that: $\forall x$ (S is speaking to $x \rightarrow x$ is F)

Intuitively, this picks out the class of properties which contains every property F which is such that S believes herself to be speaking only to Fs. One might still quibble about utterances to large audiences, in which a speaker believes that most, or almost all, of his audience has a certain property. I am not going to worry about this sort of thing.

successful. But this is very implausible; surely our verdict about the case of the beach should not be changed by simply adding in this belief on the part of the speaker.

The same argument goes the other way. Perhaps I am by nature terribly lacking in self-confidence and always feeling quite misunderstood, and that as a result my every use of a demonstrative expression, no matter how careful, is attended by a pathological belief that my audience will not grasp my intentions. If our present theory were correct, then this would entail that no demonstratives out of my mouth ever have a semantic value. But this cannot be right; surely one might correctly re-assure such a self-doubting speaker that he has, contra his worries, been using demonstratives perfectly correctly all along. Our present theory, by contrast, makes his self-doubt implausibly self-fulfilling.

Hence we need some restriction on the beliefs of the speaker which can play this role in determining the X-properties. One natural suggestion is that the only beliefs which play this role are the speaker's beliefs about the common ground. On this view, the X-properties are all those properties that the speaker believes the common ground to attribute to the audience—all those properties, that is, which the speaker believes that all the members of the conversation mutually believe to be accepted by all members of the conversation to hold of the audience of the utterance.

But it is not clear that this really does much, since we can re-create each of the above problems for the new account. (I might so lack self-confidence that I not only believe that my audience will fail to discern my referential intention, but also believe that it is part of the common ground of the conversation that this is the case; this still should not suffice for reference failure.) And it still leaves the class of beliefs which can go into determining the world relevant for evaluating the semantic value of the demonstrative unacceptably broad. The worlds determined by the properties the speaker believes the common ground to attribute to the audience might not provide a good test for the actual referential success of the speaker's utterance.

Indeed, in the worst sort of case, they will determine no world at all, for the speaker may believe that the common ground of the conversation attributes a property to the audience which, it turns out, no audience could have. Consider the following case:

Bob and Jim:

I am a student in a section of Intro to Philosophy, and I notice that students who go by the names of 'Bob' and 'Jim' never seem

to come to class on the same day. On the basis of this and other evidence, I come to form the false belief that Bob = Jim. Indeed, I take myself to know that Bob = Jim; and I take you, with whom I have discussed my evidence, to also know this, and know that I know it. We remark about how Bob/Jim, when answering to 'Bob', sits in the front row wearing jeans, and when answering to 'Jim', sits in the last row wearing a track suit. One day Bob's seat is empty, and the professor walks in and remarks on Bob's absence. Under my breath, I say to you, 'Look at him—he's just sitting in the back row like this has nothing to do with him'. Jim is in the back row along wearing a track suit along with four other male students, but I do not bother to point at him; given what I take to be our mutual knowledge of the Bob/Jim identity theory, I have good reason to believe that, in this context, it will be obvious to you that I intend to refer to Jim. And it seems that I do succeed in securing Jim as the semantic value of 'him'.

Here the speaker takes the common ground to attribute to his audience the property of knowing that Bob = Jim—but, since it is impossible to know a necessary falsehood, this is a property that no one could have. This threatens to make (Coordination-C) vacuous in any such context. (Note further that this is not just a case in which the speaker takes the common ground to attribute to his audience a property which turns out to be impossible; it is also a case in which the common ground really does attribute this property to the audience. Hence it is a problem for definitions of the X-properties in terms of the common ground just as much as for definitions of the X-properties in terms of speaker beliefs about the common ground.)

There are obvious defensive moves to be made here.

- (i) One might grant the point, but deny that this makes (Coordination-C) vacuous, on the grounds that counterpossibles are not all vacuously true.
- (ii) One might utilize the apparatus of two-dimensional semantics, and say that what determines the X-properties are the primary rather than secondary intensions of the speaker's beliefs about what the common ground attributes to the audience.

- (iii) One might simply add a brute force stipulation to the effect that any impossible properties be excluded from the X-properties.²⁴

I am not convinced that any of these defensive moves work on their own terms, but I am not going to argue that at length here. Instead I will briefly give two reasons for thinking that none of these sorts of responses solve the underlying problem.

First, suppose that I am lecturing to my Intro to Philosophy class, and believe (truly, I suspect) that the common ground of the conversation attributes to the audience the property of being inattentive. Then, given that the present version of (Coordination-C) has us look at a world in which the audience has the X-properties and is, among other things, attentive, to determine whether my actual uses of demonstratives are successful, we will have to ask whether my audience discerns my referential intentions in a world in which that audience is both attentive and inattentive. One can hardly hold that this world is, while metaphysically impossible, nonetheless epistemically possible, so it is not clear that primary intensions will be any help; the brute force exclusion will not rule out the relevant belief; and even if we are prepared to say that counterpossibles with necessary falsehoods in the antecedent are non-vacuous, it seems pretty odd to say

²⁴ One way to carry out this last strategy which is applicable to the above example would be to say that we should ignore the fact that the common ground attributes to the audience the property of knowing that Bob = Jim, and focus on the fact that it attributes to the audience the property of believing that Bob = Jim, which of course is possibly instantiated. But this particular move highlights the perils of this general strategy.

As Williamson (2000) has emphasized, there are plenty of cases in which it would be rational to take the conditional probability of an agent performing some action given his believing *P* to differ from the conditional probability of his performing that action given his knowing *P*. A special case of this generalization is that there will be occasions on which it is rational to take the conditional probability of someone taking *o* to be the value of my use of *d* in *c* given their knowing *P* to be significantly greater than that person taking *o* to be the value of my use of *d* in *c* given their merely believing *P*. But this, given the basic claim of the Coordination Account that a use of a demonstrative has a semantic value only if the speaker does something which puts the audience in a position to discern that semantic value as the intended reference, means that whether a speaker believes that an audience knows rather than merely believes a proposition can make a difference to the question of whether the demonstrative has a semantic value. Hence the move from what the common ground says about what the audience knows to what the common ground says about what the audience believes is far from innocent.

One could of course replace 'believing *P*' with 'justifiably believing *P*' or some more complicated but still non-factive condition. Williamson argues plausibly that no such condition will be an effective surrogate for knowledge. And, crucially, the more complicated the relevant condition gets, the less plausible it is to say that speakers have beliefs about whether their audience satisfies that condition.

that my actual referential success is hostage to goings on in logically impossible worlds.

Second, and relatedly: many of the problems caused by attribution of impossible properties are also caused by attribution of properties which are such that they could only be instantiated by worlds quite different from the actual world. In that case the relevant counterfactual will not take us to an impossible world; but it will take us to quite a 'distant' possible world, and the view that actual referential success or failure is to be analysed in terms of the beliefs of speakers in such possible worlds does not (as noted above) seem very attractive. This suggests, I think, that there is something fundamentally wrong with the idea that we can analyse actual referential success in terms of audience uptake in worlds where the relevant beliefs of the speaker are true.²⁵

We might continue looking for better candidates for the X-properties. But the role they are supposed to play in handling cases like sudden blindness seems to ensure that they will have to be defined in terms of some non-factive mental states of speakers, or conversational participants more generally. And this suggests that they will always lead to problems like the ones just discussed, in which the X-properties determine a world which is not plausibly relevant to actual referential success. So I think that the project of giving an adequate account of the X-properties looks, if not hopeless, pretty difficult.

Let us for a moment take a step back from the counterexamples, and think about the shape of the theory at which we have arrived. At first glance, the Coordination Account sounds like — as its name suggests — an account which explains successful use of a demonstrative partly in terms of something the speaker does, and partly in terms of something the audience does, or could do. What the speaker intends and does must be coordinated with the abilities of the audience.

But that is not the kind of theory we have ended up with. Each revision of the theory has shifted our emphasis from properties of the audience of a given utterance to properties of the mental states of the speaker. Of course, the relevant mental states are mental states which are about the audience — but, given that the relevant mental states are non-factive, this entails nothing at all about the actual properties of

²⁵ One might wonder whether (as Bryan Pickel suggested in conversation) we could get around these sorts of problems by formulating the coordination condition as a claim about the dispositions of audience members, where these are not to be analysed in terms of the obvious associated counterfactuals. Perhaps some version of this sort of view could work, though I do not see quite how it would go.

the audience. In this sense, clause (Coordination-C) under our latter interpretations is not really a condition on audiences. Consideration of the cases of mistakes about the existence and properties of audiences discussed above has led us to a theory in which what King calls the ‘hearer fact’ is not a fact about the hearer at all, but rather a further fact about the speaker: it is just a further condition on her mental states and, in particular, on her beliefs.²⁶

But the attempt to give weight to those beliefs by asking what would be the case in worlds in which they were true has led to difficulties. The question thus arises whether we can solve the problem of insufficient intentions with a suitable second condition on successful reference, involving the mental states of speakers with respect to their (believed) audience, which does not take this detour through possible worlds in which various of the speaker’s beliefs are true.

²⁶ It is a slight overstatement to say that (Coordination-C) expresses just a condition on the beliefs of the speaker; the audience does play some role. That is because the world determined by the antecedent of the embedded counterfactual in (Coordination-C) is partly determined by the beliefs of the speaker, and partly determined by facts about the actual context *c*, which of course includes the audience at *c*—it is the world-state most similar to *c* at which the relevant beliefs are true. But that the audience gets even this limited role is at the mercy of the speaker, since, in the theory we have been developing, wherever the speaker has a false belief about his audience, what matters for evaluating the success of the relevant use of the demonstrative is what the speaker believes, and not how the audience is. And, more importantly, it is not at all clear that facts about the actual context should play even the limited role in determining the success or failure of a given demonstrative utterance given them by (Coordination-C). Consider the following pair of cases:

The wine glasses:

(Case 1) A speaker is undergoing a complex hallucination in which he takes himself to be addressing someone in the room, and says, pointing to one of several wine glasses on the table, ‘That glass...’ The speaker does not have very many beliefs about his audience. As it turns out, the speaker is the only one in the room; but Jane is just outside the door. Were Jane in the room, she would have been able to tell which wine glass the speaker was referring to.

(Case 2) As above, except that Sally is outside the door. Sally has an unusual phobia regarding wine glasses, and, upon catching a glimpse of the wine glasses on the table, would immediately look away. So Sally would not have been in a position to tell which wine glass the speaker is demonstrating.

It seems that, on (Coordination-C), the difference between these cases matters. The nearest world at which the antecedent of (Coordination-C) is true in Case 1 will be one in which Jane is in the room, and analogously for Case 2 and Sally. Hence in Case 1 ‘that glass’ will have a semantic value, and in Case 2 it will not. This seems bizarre; surely we should say the same thing about these two possible utterances of ‘that glass’. The psychological properties of the person outside the room—which might of course be unknown to the speaker—just should not matter. This suggests that letting the actual properties of audience members play even the limited role given them by (Coordination-C) is a mistake.

5. Rejecting the internalist intuitions

In the next section, I will develop an account along these lines. But first I want to briefly consider a different line of response from the defender of the Coordination Account. Perhaps the best move for the proponent of that account is not to continue further down the road we are on, but to reject some of the intuitions that got us started on that road in the first place. Could we, for example, just deny the intuition that ‘that lectern’ has a semantic value in the example of the sneaky students?

One might think that this could be motivated by analogy with another topic in the theory of reference. Recall Russell’s discussion in ‘Knowledge by Acquaintance and Knowledge by Description’ of the scope of our acquaintance, and the subsequent view that the only constituents of propositions which we can grasp must be sense data, properties, relations, and (perhaps) oneself. Why did Russell hold this restrictive view? On one plausible reading, he was moved by something like the following sort of argument:

Whatever I can be mistaken about, I cannot be mistaken about whether I am thinking a thought. But suppose that I could introduce a genuine name *n*—a name whose meaning was its reference—for an object *o* about whose existence I could be mistaken. Then I could be mistaken about whether I was thinking a thought because, by using using a sentence involving *n*, there would be no thought that I was thinking. Hence I cannot name anything about whose existence I can be mistaken.

Russell’s conclusion is now widely regarded as false. But if it is false, then something must be wrong with his argument; and if something is wrong with his argument, it must be the case that I can sometimes use an expression in what seems to me a perfectly good way, but that, just in virtue of the world failing to cooperate, my use of that expression lacks a semantic value (and hence sentences involving it fail to express a thought).

The defender of the Coordination Account might then reasonably ask whether the intuitions which guide us in the case of the sneaky students or the timid burglar or sudden blindness are coming from the same discredited source that fed Russell’s thought that it is impossible to name something about whose existence one can be mistaken. Once we admit that whether our expressions have semantic values is not completely ‘up to us’, but also involves the cooperation of our

environment, why not also say that whether ‘that lectern’ has a semantic value in the context depends on whether his audience stuck around to hear it? Why would the case of the sneaky students not be, just as much as a hallucination, a way of the world failing to cooperate in the way needed for a demonstrative to have a semantic value?

If one is looking for a way to block the gradual but, once started pretty inevitable, slide to the various versions of (Coordination-C) before it starts, it seems that we need to start with a condition even stronger than (Coordination-A). For (Condition-A) permits referential success in cases in which the audience does not discern the speaker’s referential intention, but would have if the audience were competent and attentive. But why should this be permitted? Presumably because my audience’s failing to pay attention is not fault of mine, and cannot stop me from successfully using a demonstrative.²⁷ But this intuition—that things which, through no fault of my own, stop my audience from discerning my intentions cannot by themselves bring about reference failure—is just what fuels the intuitions about cases that lead to the C-versions of the Coordination Account.

Better, if this slide is to be avoided, to hold that referential success requires one’s audience not just to be such that they *would* discern one’s referential intentions if competent and attentive, but actually to discern one’s referential intentions. This, I think, is what the proponent of the view that the character of demonstratives is partly to be given in terms of ‘hearer facts’ should say; it draws a principled distinction between referential failure and success, and seems to me much more promising than to attempt to solve the problem to which speaker beliefs give rise discussed in the last section.²⁸

But the view has some serious costs. One is simply that it is committed to quite surprising claims about the truth conditions of sentences involving demonstratives. Think about the following two cases:

²⁷ As King puts it: ‘we don’t want to require that the hearers *in fact* figure out what the relevant value is. That would be too strong. If my hearer is inattentive, incompetent or simply ignoring me, that should not by itself prevent me from securing a value for my demonstrative’ (2014b, p. 225).

²⁸ Mount (2008a) makes claims which suggest that she thinks that actual audience uptake is required for referential success: ‘An object is *mutually recognized as maximally salient* by conversational participants when all interlocutors have focused their attention on it, and are aware that they have all focused their attention on it. Thus my claim is that a demonstrative refers to the object mutually recognized as maximally salient. When there is no such object, the demonstrative does not refer’ (pp. 154–5).

case 1 is just the case of the sneaky students; in case 2, one student fell asleep even before the lights went down, and hence failed to leave the lecture room with his compatriots, and woke up just in time to hear ‘that lectern’. We are contemplating a view on which the presence of this single remaining student makes all the difference: if the student is there in the 14th row, then ‘that lectern’ has a semantic value; if the student leaves with the rest of them, then it lacks one. This seems to me a very big bullet to bite.²⁹

Perhaps not all will agree that this result is so problematic. But there is also a more basic puzzle about this sort of radical externalist view. We do not ordinarily take the semantic value of an expression in a context to coincide with the audience’s view of the semantic value of that expression. Misunderstanding is possible. But then it is a bit puzzling that we should say that these things must coincide just in the case of the demonstratives. Why should demonstratives, lone among expressions, be such that it is metaphysically impossible to use them with a semantic value unless that semantic value is grasped by one’s audience?

So rather than multiplying examples which are relevantly similar to the one just given, I will simply note this as one way in which the proponent of the Coordination Account might try to block the main line of argument to this point, and turn to the different question of how we might solve the problem of insufficient intentions without moving to a theory which so radically disregards intuitions about the truth conditions of the target sentences.

6. The problem of insufficient intentions and modified intention theories

In this section, I want to do two things: express some scepticism about the idea that the problem of insufficient intentions is really a problem for the Simple Intention Theory; and then say how, if it is a problem, it can be solved.

One might resist the idea that the problem of insufficient intentions really is a problem by resisting King’s claim that cases like the example of the beach (or the similar cases discussed by Wettstein,

²⁹ Related problems arise in cases where a speaker is addressing a large audience. We presumably do not want to say that, for ‘that lectern’ to have a semantic value, each of the 235 students in Intro to Philosophy must recognize the object of my referential intention. But then how many are required? It is hard to imagine coming up with any principled answer to this question.

Reimer, and Gauker) really are cases in which the demonstrative lacks a semantic value. The intuition that they do lack a semantic value is due to the fact that the speaker in those examples fails to put his audience in a position to discern her referential intention — and, in both cases, would be able to see with just a bit of reflection that she will fail in this way. The problem is that there are commonplace uses of demonstratives — which, in this context, were brought to my attention by Eliot Michaelson — which are like the example of the beach in this way, but seem to be clear cases of referential success.

Imagine that we are sitting on a couch in my house, and I think I see something quite surprising — like a bird quickly flying past the doorway. You do not flinch, so I am almost sure that you did not see what I think I saw. But, to be sure, I might ask you: ‘Did you see that?’ Supposing that a bird really did fly past the doorway, it seems clear that my use of ‘that’ here succeeds in referring to the bird. But did I do anything to make that referential intention clear to you? It seems pretty clear that I did not. And, indeed, we can suppose that this fact is clear to me, or would have been after a moment’s reflection — after all, I do not suppose you to be so implacable as to just sit there while you saw a bird fly past. None of these suppositions makes it any less plausible that my demonstrative had the bird as its semantic value.

One might, of course, just press cases like this as a counterexample to the various versions of the Coordination Account — but here I want to ask, instead, whether this sort of case casts doubt on the verdict of reference failure in the examples used to generate the problem of insufficient intentions. The challenge posed by this sort of case is just for the proponent of the verdict of reference failure in the case of the beach to provide some sort of principled distinction between the two cases. The cases are, after all, alike in some important ways: both are cases in which the audience is in no position to determine the value the speaker intended the demonstrative to have; and both are cases in which the speaker knows this, or could easily know it upon reflection. So it looks like the relevant ‘hearer facts’ are the same, and the proponent of the Coordination Account will be forced to treat the two cases as parallel.

But then the proponent of the Coordination Account will have to say either that the beach is not a case of reference failure, or that the example of the bird is a case of reference failure. Since the latter does not seem especially plausible, we seem forced into saying that the beach is not, after all, a case of reference failure. But then the problem

of insufficient intentions dissolves, and with it the counterexample to the Simple Intention Theory, and the corresponding motivation for the Coordination Account.

One might, of course, try to find some relevant difference between the cases. One which stands out is that, in the case of the beach, the speaker could have easily done more to make his referential intention clear, whereas I had no such obvious recourse in the example of the bird flying by the doorway. So the speaker in the case of the beach is irresponsible in a certain way that the speaker in the case of the bird is not. This is true, but seems to be an implausible place to draw the line between referential failure and success; one typically does not think that irresponsibility on the part of a speaker should suffice for the words out of her mouth to lack a semantic value. Hence it is at least reasonable to wonder whether our sense that something has gone wrong in the case of the beach is tracking this sort of negative judgement about the speaker, rather than anything about the semantic values of the demonstratives or truth conditions of the utterance.³⁰

But this is not a point that I want to press any further. Even if for the sake of argument we agree with King and others that cases like the beach really are cases of reference failure, we have seen that the resulting problem of insufficient intentions is not one to which the Coordination Account gives an adequate solution. It remains to be shown that this problem can be solved by a view which explains the character of demonstratives wholly in terms of the properties of the speaker of the context.

And in fact to solve the problem here we need not invoke any facts about the audience; we can simply strengthen the Simple Intention Theory by putting some extra conditions on the mental states of the speaker. Here are two suggestions:

Intention + Intention (I+I):

The value of a use of a demonstrative *d* in a context *c* is *o* iff:

- (1) the speaker intends *o* to be the value of *d* in *c*

³⁰ Robert Stalnaker suggested to me another way of arguing for this conclusion, which parallels one of the central arguments of Kripke 1977. We can imagine a language in which demonstratives are stipulated to have the character assigned to them by the Simple Intention Theory. But even in a community using this language, we can imagine utterances like the one in the case of the beach; and speakers of that language would, like us, have the intuition that something was wrong with these utterances. This casts some doubt on the explanation of that intuition as tracking reference failure.

- (2) the speaker intends that his audience take *o* to be the object that the speaker intends to be the value

Intention + Justified Belief (I+JB):

The value of a use of a demonstrative *d* in a context *c* is *o* iff:

- (1) the speaker intends *o* to be the value of *d* in *c*
- (2) the speaker would be justified in believing that his audience will take *o* to be the object that the speaker intends to be the value³¹

These are not unrelated to the Coordination Account.³² Both are like that account in focusing, not just on intentions to refer, but on intentions and/or beliefs about an audience.

This gives each some of the virtues of the Coordination Account. Just as the Coordination Account can say that demonstrations are sometimes required, and sometimes not, just because they are only sometimes necessary for an audience to figure out what object the speaker intends to refer to, I + I can say that demonstrations are sometimes (but not always) a necessary means to carrying out a speaker's audience-directed intention, and I + JB can say that demonstrations are sometimes (but not always) necessary for the speaker to be justified in the belief that his audience will be in a position to recognize which object he intends to be the value of the demonstrative. Both also, like the Coordination Account, incorporate reference to an audience while remaining smoothly generalizable to the case of talking to oneself. It is just that neither cashes this 'reference to an audience' out in terms of what some audience does or would do. For this reason, neither of these theories gets into the troubles discussed in section 4 above.

Do these theories give us a verdict of reference failure in the case of the beach? I + JB, at least, seems to. The speaker is not justified in

³¹ Note that this should not be understood as requiring the speaker to have the relevant belief—otherwise we would get into the same trouble as did (Coordination-C) with speakers who so lack self-confidence in their use of demonstratives that they never expect to be understood.

Here as elsewhere I am setting aside Kaplan's idea that character should be not just what determines content relative to a context, but also what competent users of the expression must know about its meaning. One might reasonably doubt whether competent users of demonstratives must have anything like I + I or I + JB in mind. Of course, the same concern would apply *a fortiori* to the Coordination Account, especially in its more complex forms.

³² I + I is also closely related to the view defended in Bach 1992.

believing that his audience will figure out which woman he is talking about, and this entails that his use of the demonstrative lacks a value. Moreover, if the speaker were so justified—for example, if he were justified in believing that his interlocutors were mind-readers or preternaturally good at interpreting vague waves of his hand—then our intuition of reference failure vanishes. This is exactly what I + JB predicts.

Matters are less clear with I + I. But I think that a reasonable case can be made that when we think of the case of the beach as involving reference failure, we are imagining the speaker sort of carelessly waving his hand in the direction of the woman to whom he intends to refer, without caring whether he is understood. In that sort of case, he also lacks the second intention required by I + I.³³

Given that I can intend to do something without being justified in believing that I will do it—as when I intend to make a half-court shot in basketball—these theories are not equivalent. (This fact might count in favour of I + I, since just for this reason it plausibly gives the right result in the example of the bird flying past the doorway, whereas I + JB does not.³⁴) I am not sure which theory is better—or indeed whether their conjunction, I + I + JB, is the way to go.

Here, I am not going to try to decide this question. I think that—at least if you share my intuitions about the cases discussed above—it is clear that any of these options is preferable to any version of the Coordination Account we have discussed. They handle, and more simply, all of the cases that the Coordination Account handles, and run into none of the problems for the Coordination Account detailed in section 4. To return to the big picture question sketched at the outset—about whether we should supplement intention-based approaches to the character of indexicals with further facts about the psychology of the speaker or with facts about the audience—this shows that, so far as the problem of insufficient intentions goes, theories of the former sort are to be preferred.

³³ I + I also seems to me to say the right thing about Gauker's example of the tie. It seems to me that there are two readings of that case: one in which the speaker is intending to address the tie-wearer, and one in which she is really just thinking out loud, and talking to herself. The latter case is unproblematic for reasons already discussed in connection with the case of the idiot. And to me, the former case no longer seems like a case of reference failure if we stipulate that the speaker really did intend that her audience recognize the object of her referential intention.

³⁴ Thanks to Robin Jeshion for pointing this out.

But that is not the final word. The Simple Intention Theory, after all, faced difficulties not just with the problem of insufficient intentions, but also with the problem of conflicting intentions. And this is true no less of the modified psychological theories just sketched. All we need to do is imagine a version of the case of Carnap and Agnew in which the speaker intends his audience to see that he is pointing at a picture of Carnap, and is justified in believing that they will do so, and we get a counterexample to each of the Simple Intention Theory, I + I, and I + JB. Hence one might think that, despite the problems detailed above, the Coordination Account might still receive a powerful motivation from the problem of conflicting intentions.

7. Coordination and the problem of conflicting intentions

King himself, in ‘Speaker Intentions in Context’ (King 2014b), takes the demonstrative in the case of Carnap and Agnew to lack a semantic value, and hence does not think of the Coordination Account as having any special advantage in solving the problem of conflicting intentions.³⁵ But, as noted above, one might reasonably think that King has underestimated the advantages of his own theory here. It does after all seem quite plausible that the demonstrative in that example successfully refers to the picture of Agnew hanging on the wall, and one might, as noted above, reasonably think that the Coordination Account, unlike the Simple Intention Theory, delivers this result.

The reason was simply that the picture at which the speaker is pointing, but not the picture of Carnap which used to hang on the wall, satisfies the Coordination Condition. A competent, attentive, reasonable audience would take the speaker to be referring to the picture at which the speaker is pointing, and would not take the speaker to be referring to the picture of Carnap. Hence we seem to get, without need for any extra machinery, a solution to our problem.

Unfortunately, there are two problems lurking. The first is that the above line of reasoning, while plausible enough if we have in mind the A- or B-versions of the coordination condition, gets no purchase once we work our way to (Coordination-C). For in the Carnap/Agnew case, the speaker believes that the common ground attributes to the audience the property of seeing above his head a picture of Carnap. And of

³⁵ He expresses some second thoughts about this view, and develops an alternative account, in King 2013.

course in the nearest world in which that belief is true (and the audience is attentive, competent, etc.) the audience will discern the picture of Carnap, rather than the picture of Agnew, as the object of his referential intention. Hence the proponent of the C-version of the Coordination Account can offer no treatment of Carnap and Agnew. But, as we saw, there is very strong reason for the proponent of the Coordination Account to adopt (Coordination-C)—for otherwise we get quite counterintuitive results in cases in which the speaker has false beliefs about the properties of her audience.

The second problem is that there are cases of conflicting intentions which no version of the Coordination Account can handle. These are cases in which both the speaker and the audience have false beliefs about their environment. Here is such a case:

Mutual hallucination:

Suppose that you and I are having a conversation while jointly experiencing a quite convincing mutual hallucination of your cat ‘Fluffy’. Pointing at what I take to be the Fluffy, I say ‘That cat ...’

This may not at first seem like a case of conflicting intentions, but it is. I intend to refer to Fluffy, but also intend to refer to the cat at which I am pointing. The first intention singles out Fluffy; the second intention singles out nothing. The fact that this is a case of reference failure is simply due to the fact that the second intention trumps the first.

But the Coordination Account can give no account of this fact. Fluffy, after all, satisfies both conditions on reference: she is the object of my referential intention, and you are readily able to discern, given our mutual hallucination, that she is the object of my referential intention.

8. The unsolved problem of conflicting intentions

It would be pleasingly symmetrical to end this paper with a purely speaker-based solution to the problem of conflicting intentions in the spirit of the solutions to the problem of insufficient intentions offered above. As the title of this section suggests, I am not in a position to give such a solution. Nor, I think, is anyone—in this respect, the proponent of the Coordination Account is no worse off than anyone else. The aim of this section is to say something about the obstacles that stand in the way of a solution to this problem.

It would be nice to be able to avoid this problem altogether by saying, as King (2014b) does, that every case of conflicting intentions is a case of reference failure.³⁶ But, while attractively simple, this line of response looks considerably less plausible when we consider just how common cases of conflicting intentions are. Imagine a carnival where, for a fee, I can choose from among a large number of plastic balls, one of which contains \$100. The person running the game asks me which one I want and, pointing, I say, ‘That ball’. I intend to refer to the ball at which I am pointing (say, ball #58), and of course also intend to refer to the ball which contains \$100. But it turns out that another ball — ball #113 — contains the cash. The intentions conflict — but it is clear that the semantic value of my utterance of ‘that ball’ is ball #58.

This suggests that any successful treatment of demonstratives must provide some account of the conditions under which one intention to refer trumps another. A natural way to execute this strategy is to try to find some special intrinsic quality of intentions which makes them always trump other referential intentions. The problem is that it seems that we can always find cases in which we have conflicting intentions, both with the alleged special intrinsic quality.

For example, when thinking about the case of the carnival just described, one might think that *de re* intentions always trump. But we can imagine a variant of the above case in which the speaker has a prior acquaintance with the winning ball, and a *de re* intention to refer to it. We would then have conflicting *de re* intentions. (And, even worse, in the original Carnap/Agnew case, we seem to have a descriptive intention trumping a *de re* intention.)

But perhaps the category of *de re* intentions is simply the wrong choice for the class of always-trumping intentions. One notable feature of the cases of conflicting intentions discussed so far is that demonstrations seem to play some special role; intentions to refer to the thing one is demonstrating seem always to trump other sorts of referential intentions. Hence, adapting the view of Reimer (1992), herself adapting the view of Kaplan (1979), we might hold that the intentions which accompany demonstrations play a special role here; perhaps the trumping intention is always, in Reimer’s sense, a secondary intention. In the case of Carnap and Agnew, this would be ‘an intention to demonstrate, and say something of, the picture in the direction of the gesture’ (Reimer 1992, p. 390). More generally, on this view, the

³⁶ As noted above, King is now less inclined to take this line. See King 2013.

trumping intention will be, for some property *F*, the intention to refer to the thing which instantiates *F* in the direction of the gesture.

But this proposal, while initially attractive, faces a number of problems. First, there are cases in which a speaker intends to refer to the *F* in the direction of her pointing gesture, and succeeds in referring to *o*, despite the fact that *o* is not *F*. I might say when teaching a lecture class, 'Wake that student up', pointing at a student in the front row slumped on his desk with his eyes closed, and intending to refer to the sleeping student in the direction of my pointing gesture. 'That student' might well have the apparently sleeping student as its semantic value, even if the student is protesting the tedium of the lecture by pretending to sleep.³⁷

In other cases, I can have conflicting intentions, both of which are, in Reimer's sense, secondary intentions. That is, I might intend to refer to the *F* in the direction of my pointing gesture and intend to refer to the *G* in the direction of my pointing gesture, in a situation in which the *F* $\neq$ the *G*. To generate such a case, just imagine that next to the student pretending to be sleeping, there is another student sleeping with eyes wide open. I intend to refer to the sleeping student in the direction of my pointing gesture, and also intend to refer to the student with eyes shut in the direction of my pointing gesture. This will be a case of conflicting intentions which are secondary, in Reimer's sense. And if we can have conflicting secondary intentions, we cannot solve the problem of conflicting intentions by claiming that secondary intentions always trump.

Finally, there are cases of conflicting intentions which involve no demonstration at all.³⁸ One is a simple variant of the example of Carnap and Agnew. Suppose that there is initially no picture on the wall behind the speaker, but that some workmen noisily wheel out a large canvas behind the speaker. The plan was for them to wheel out the picture of Carnap; but they have mistakenly wheeled out a picture of Spiro Agnew. When the speaker says 'That is a picture of one of the greatest philosophers of the twentieth century', he intends to refer to the picture of Carnap, and of course also intends to refer to the picture noisily wheeled out by the workmen just before his utterance. If one

³⁷ One might try to solve this problem by making the value of '*F*' something very general, like 'thing.' But this would reinstate one of the central problems of the pure demonstration theory of demonstratives, which is that in pointing at one thing, I am always also pointing at many other things.

³⁸ Thanks to an anonymous referee for this point.

thinks that, in the original version of the case, the demonstrative refers to the picture of Agnew, it is hard to resist saying the same about this variant on the case — despite the absence of any demonstration by the speaker.

So far our search for the class of always-trumping intentions has come up empty. A more promising strategy for solving the problem of conflicting intentions gives up this search, and instead tries to analyse trumping in terms of explanatory relations between intentions in the relevant context of utterance.³⁹ But it is at least not easy to see how to give an account of this sort, since explanatorily prior relations sometimes are trumped, and other times do the trumping.

In the case of the carnival, I intend my demonstrative to refer to ball #58 because I intend it to refer to the winner — and the latter, explanatorily prior, intention, gets trumped. So this suggests that the intention which comes last in the causal chain in a particular context should do the trumping.

But now recall the case of the student who seems to be, but is not, sleeping. I might see him with his eyes closed, and form the *de re* intention to refer to that student. Because I have this *de re* intention, I might point in his direction and form the intention to refer to the sleeping student in the vicinity of my pointing gesture. The latter intention is explanatorily posterior, so we should expect it to do the trumping; but it does not, and instead seems to be trumped by the explanatorily prior *de re* intention to refer to that student.

So, while it seems like the problem of conflicting intentions must have some solution, and it seems quite plausible that it will have to be based on explanatory or causal relations between intentions in the context of utterance, it is not at all easy to see how to turn this general thought into an adequate theory.⁴⁰

9. Conclusion

We began with the idea that various sorts of context-sensitive expressions might best be handled by making their characters sensitive to the intentions of the speaker of the context. We then considered various

³⁹ Promising accounts of this sort have been developed by Perry (2009), who develops the view of Kaplan (1989), and by King (2013). The idea that semantic properties are fixed by intentions which play a certain explanatory role in the psychology of the speaker has been explored in many places; see, for example, Davidson 1986.

⁴⁰ I discuss this problem in more depth in Speaks MS.

problems for this approach, and considered two ways of solving these problems: King's route of making those characters sensitive to various 'hearer facts' — properties of the audience of the context — and the more conservative route of making those characters sensitive, not just to the referential intentions of speakers, but also to various other aspects of speaker psychology. We have now, I think, found some very strong reasons to prefer the latter route.

First, we have failed to identify any problem for the Simple Intention Theory which is solved by the Coordination Account and is not solved more simply by a purely speaker-based theory. Second, in response to counterexamples, we have found that, under pressure, the Coordination Account collapses into an account which relies, not on properties of audiences, but properties of beliefs of speakers about audiences — and hence into a type of purely speaker-based theory. Third, the type of purely speaker-based theory into which it collapses — one which makes use of some of Coordination-C — itself leads to unacceptable consequences, as illustrated by the examples of speakers with inconvenient beliefs about their audiences discussed at the end of section 4.

But the discussion of the problem of conflicting intentions should give proponents of speaker-intention-based approaches to context-sensitivity some pause. This is not because a focus on the audience, as the Coordination Account would recommend, fares better — as we saw in section 7, it does not. Rather, it is just because, if it really is true that the semantic values of indexicals are often fixed by the intentions of speakers, we should be able to find some principled way of saying how conflicts between those intentions are resolved. As I have tried to show in this section, that is easier said than done.

So, while I think that the foregoing gives us good reason to prefer a purely speaker-based account of the character of demonstratives over King's, proponents of the former sort of account still have some work to do. The problem of conflicting intentions is a hard problem — the hardest, I think, for any intention-based approach to the character of demonstratives in particular, and indexical expressions more generally.⁴¹

⁴¹ Thanks for helpful discussion to Joshua Armstrong, Kent Bach, David Chalmers, John Hawthorne, Robin Jeshion, Jeff King, Eliot Michaelson, Robert Stalnaker, Juhani Yli-Vakkuri, and discussions at the 4th Semantic Content Workshop at the University of Barcelona, a conference on Reference and Frege Puzzles at the University of Umea, and a conference in Internalism in Semantics at McGill University.

References

- Almog, J., J. Perry, and H. Wettstein (eds) 1989: *Themes From Kaplan*. Oxford: Oxford University Press.
- Almog, Joseph and Paolo Leonardi (eds) 2009: *The Philosophy of David Kaplan*. Oxford: Oxford University Press.
- Burgess A. and B. Sherman (eds) 2014: *New Essays in Metasemantics*. Oxford: Oxford University Press.
- Bach, Kent 1992: 'Intentions and Demonstrations'. *Analysis*, 52, pp. 140–6.
- Chomsky, Noam 1976: *Reflections on Language*. Temple Smith.
- Cohen, Jonathan and Eliot Michaelson 2013: 'Indexicality and the Answering Machine Paradox'. *Philosophy Compass*, 8, pp. 580–92.
- Davidson, Donald 1986: 'A Nice Derangement of Epitaphs'. In Lepore 1986.
- French, Peter A., Theodore E. Uehling Jr, and Howard K. Wettstein (eds) 1977: *Studies in the Philosophy of Language*. Minneapolis: University of Minnesota Press.
- Gauker, Christopher 2008: 'Zero Tolerance for Pragmatics'. *Synthese*, 165, pp. 359–71.
- Glanzberg, Michael MS: 'Not All Contextual Parameters Are Alike'.
- Harman, Gilbert 1974: 'Review of Stephen Schiffer, *Meaning*'. *Journal of Philosophy*, 71, pp. 224–9.
- Hawthorne, John and David Manley 2012: *The Reference Book*. Oxford: Oxford University Press.
- Kaplan, David 1970: 'Dthat'. *Syntax and Semantics*, 9, pp. 221–43.
- 1977/1989: 'Demonstratives'. In Almog, Perry, and Wettstein 1989.
- 1989: 'Afterthoughts'. In Almog, Perry, and Wettstein 1989.
- King, Jeffrey C. 2001: *Complex Demonstratives: A Quantificational Account*. Cambridge, MA: MIT Press.
- 2013: 'Supplementives, the Coordination Account, and Conflicting Intentions'. *Philosophical Perspectives*, 27, pp. 288–311.
- 2014a: 'The Metasemantics of Contextual Sensitivity'. In Burgess and Sherman 2014, pp. 97–118.
- 2014b: 'Speaker Intentions in Context'. *Noûs*, 48, pp. 219–37.
- Kripke, Saul A. 1977: 'Speaker's Reference and Semantic Reference'. In French, Uehling, and Wettstein 1977.
- Lepore, Ernest (ed.) 1986: *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Oxford: Blackwell.

- McGinn, Colin 1981: 'The Mechanism of Reference'. *Synthese*, 49, pp. 157–86.
- Mount, Allyson 2008a: 'Intentions, Gestures, and Salience in Ordinary and Deferred Demonstrative Reference'. *Mind and Language*, 23, pp. 145–64.
- 2008b. 'The Impurity of "Pure" Indexicals'. *Philosophical Studies*, 138, pp. 193–209.
- Perry, John 1991: 'Demonstratives, Demonstrations, and Demonstrata'. *Philosophical Studies*, 63, pp. 187–202.
- 2009: 'Directing Intentions'. In Almog and Leonardi 2009.
- Reimer, Marga 1991: 'Demonstratives, Demonstrations, and Demonstrata'. *Philosophical Studies*, 63, pp. 187–202.
- 1992: 'Three Views of Demonstrative Reference'. *Synthese*, 93, pp. 373–402.
- Russell, Bertrand 1910: 'Knowledge by Acquaintance and Knowledge by Description'. *Proceedings of the Aristotelian Society*, 11, pp. 108–28.
- Speaks, Jeff 2008: 'Conversational Implicature, Thought, and Communication'. *Mind and Language*, 23, pp. 107–22.
- MS: 'The Problem of Conflicting Intentions'.
- Stalnaker, Robert 2002: 'Common Ground'. *Linguistics and Philosophy*, 25, pp. 701–21.
- Wettstein, Howard K. 1984: 'How to Bridge the Gap Between Meaning and Reference'. *Synthese*, 58, pp. 63–84.
- Williamson, Timothy 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.